arXiv:1403.4689v1 [math.PR] 19 Mar 2014

Exponential Family Techniques for the Lognormal Left Tail

Søren Asmussen[1], Jens Ledet Jensen[1] and Leonardo Rojas-Nandayapa[2]

[1]Department of Mathematics, Aarhus University
[2]School of Mathematics and Physics, University of Queensland

Abstract

Let X be lognormal(μ, σ^2) with density $f(x)$, let $\theta > 0$ and define $\mathcal{L}(\theta) = \mathbb{E}\mathrm{e}^{-\theta X}$. We study properties of the exponentially tilted density (Esscher transform) $f_\theta(x) = \mathrm{e}^{-\theta x} f(x)/\mathcal{L}(\theta)$, in particular its moments, its asymptotic form as $\theta \to \infty$ and asymptotics for the Cramér function; the asymptotic formulas involve the Lambert W function. This is used to provide two different numerical methods for evaluating the left tail probability of lognormal sum $S_n = X_1 + \cdots + X_n$: a saddlepoint approximation and an exponential twisting importance sampling estimator. For the latter we demonstrate the asymptotic consistency by proving logarithmic efficiency in terms of the mean square error. Numerical examples for the c.d.f. $F_n(x)$ and the p.d.f. $f_n(x)$ of S_n are given in a range of values of σ^2, n, x motivated from portfolio Value-at-Risk calculations.

Keywords: Lognormal distribution, Esscher transform, exponential change of measure, Laplace transform, Laplace method, Cramér function, saddlepoint approximation, Lambert W function, rare event simulation, importance sampling, VaR.

MSC: 60E05, 60E10, 90-04

1 Introduction

The lognormal distribution arises in a wide variety of disciplines such as engineering, economics, insurance or finance, and is often employed in modeling across the sciences [2, 12, 14, 20, 21]. In consequence, it is natural that sums of lognormals come up in a number of contexts. For instance, a basic example in finance is the Black-Scholes model, which asserts that security prices can be modeled as independent lognormals (equivalently, the logprices are independent normally distributed). This implies that the value of a portfolio with n securities can be conveniently modeled as a sum of lognormals. Another example occurs in the valuation of arithmetic Asian options where the payoff depends on the finite sum of correlated lognormals [13, 23]. In insurance, individual claim sizes are often modeled as independent lognormals,

so the total claim amount after certain period is a random sum of lognormals [26]. A further example occurs in telecommunications, where the inverse of the signal-to-noise ratio (a measure of performance in wireless systems) can be modeled as a sum of i.i.d. lognormals [17].

However, the distribution of a sum of n lognormals S_n is not available in explicit form and its numerical approximation is considered to be a challenging problem. In consequence, a number of methods for its evaluation has been developed across several decades, but these can rarely deliver arbitrary precisions in the whole support of the distribution, particularly in the tails. The later case is of key relevance in certain applications which often require to evaluate tail probabilities at very high precisions. For instance, the Value-at-Risk (VaR) is an important measure of market risk defined as an appropriate $(1-\alpha)$ quantile of the loss distribution, and the standard financial treatise Basel II [1] asks for calculations of the VaR for so small values as $\alpha = 0.03\%$.

When considering lognormals sums, the literature has sofar concentrated on the right tail (with the exception of the recent paper [18] by Gulisashvili & Tankov). In this paper, our object of study is rather the left tail and certain mathematical problems that naturally come up in this context. To be precise, let Y_i be normal(μ_i, σ_i^2) (we don't at the moment specify the dependence structure), let $X_i = \mathrm{e}^{Y_i}$ and $S_n = X_1 + \cdots + X_n$. We then want to compute $\mathbb{P}(S_n \leq z)$ in situations where this probability is small.

An obvious motivation for this problem comes from the VaR problem. Here S_n may represent the future value of the portfolio. If Π is the present value, $\Pi - S_n$ is then the loss, and so calculation of α-quantiles are equivalent to left tail calculations for S_n. A further example occurs in the wireless systems setting, where an outage occurs when the signal-to-noise ratio exceeds a large threshold. The outage probability is therefore obviously related to the left tail probability of a lognormal sum.

The problem of approximating the distribution of a sum of i.i.d. lognormals has as mentioned a long history. The classical approach is to approximate the distribution of a sum of i.i.d. lognormals with another lognormal distribution. This goes back at least to Fenton [15] in 1960 and it is nowadays known as the *Fenton-Wilkinson method*; according to Marlow [22] this approximation was already used by Wilkinson since 1934. However, the Fenton-Wilkinson method, being a central limit type result, can deliver rather inaccurate approximations of the distribution of the lognormal sum when the number of summand is rather small or the dispersion parameter is too high—in particular in the tail regions. Another topic which has been much studied recently is approximations and simulation algorithms for right tail probabilities $\mathbb{P}(S_n \geq y)$ under heavy-tailed assumptions and allowing for dependence, see in particular [5, 8, 10, 16, 24]. For further literature surveys, see [18].

Our approach is to use the saddlepoint approximations and a closely related simulation algorithm based on the same exponential change of measure. This requires i.i.d. assumptions, in particular $\mu_i \equiv \mu$, $\sigma_i^2 \equiv \sigma^2$. Since μ is just a scaling factor, we will assume $\mu = 0$. The saddlepoint approximation occurs in various (closely related) forms, but all involve the function $\kappa(\theta) = \log \mathcal{L}(\theta)$ where

$$\mathcal{L}(\theta) = \mathbb{E}\mathrm{e}^{-\theta X_i} = \int_0^\infty \mathrm{e}^{-\theta x} f(x)\,\mathrm{d}x \quad \text{with} \quad f(x) = \frac{1}{x\sigma\sqrt{2\pi}}\mathrm{e}^{-\log^2 x/2\sigma^2}$$

and its two first derivatives $\kappa'(\theta)$, $\kappa''(\theta)$ [note that since the right tail of the lognormal distribution is heavy, these quantities are only defined for $\theta \geq 0$]. Define the exponentially tilted density $f_\theta(x)$ (Esscher transform) by $f_\theta(x) = \mathrm{e}^{-\theta x - \kappa(\theta)} f_\sigma(x)$, and let its corresponding c.d.f./probability distribution be F_θ with expectation operator $\mathbb{E}_\theta$. Then

$$\kappa'(\theta) = -\mathbb{E}_\theta X\,, \quad \kappa''(\theta) = \mathrm{Var}_\theta\, X \tag{1.1}$$

and can connect the given distribution of S_n (corresponding to $\theta = 0$) to the $\mathbb{P}_\theta$-distribution by means of the likelihood ratio identity

$$\mathbb{P}(S_n \in A) \;=\; \mathbb{E}_\theta\big[\exp\{\theta S_n + n\kappa(\theta)\};\, S_n \in A\big]\,.$$

The details of the saddlepoint approximation involve writing $z = nx$, defining the *saddlepoint* or *Cramér function* $\theta(x)$ as the solution of the equation $\kappa'\big(\theta(x)\big) = -x$ and taking $\theta = \theta(x)$. This choice of θ means that $\mathbb{E}_\theta S_n = z$ so that the $\mathbb{P}_\theta$-distribution is centered around z and central limit expansions apply. For a short exposition of the implementation of this program in its simplest form, see [3, p. 355].

The application of saddlepoint approximations to the lognormal left tail appears first to have appeared in the third author's 2008 Dissertation [25], but in a more incomplete and preliminary form than the one presented here. A first difficulty is that $\kappa(\theta)$ is not explicitly available for the lognormal distribution. However, approximations with error rates were recently given in the companion paper [7]. The result is in terms of the Lambert W function $\mathcal{W}(a)$ [11], defined as the unique solution of $\mathcal{W}(a)\mathrm{e}^{\mathcal{W}(a)} = a$ for $a > 0$. The expression for $\kappa(\theta)$ in [7] is

$$\mathcal{L}(\,\theta) \;=\; \frac{\exp\left\{-\,\dfrac{\mathcal{W}^2(\theta\sigma^2) + 2\,\mathcal{W}(\theta\sigma^2)}{2\sigma^2}\right\}}{\sqrt{1+\mathcal{W}(\theta\sigma^2)}}\mathbb{E}\big[g_0(\sigma_\theta Z)\big], \tag{1.2}$$

where $Z \sim \mathrm{N}(0,1)$, $\sigma_\theta^2 = \sigma^2/(1+\mathcal{W}(\theta\sigma^2))$ and g_0 is a certain function such that $\mathbb{E}\big[g_0(\sigma_\theta Z)\big]$ is close to 1 (see Section 2 for more detail; we also give an extension to expectations of the form $\mathbb{E}X\mathrm{e}^{-\theta X}$ there).

The paper is organized as follows. In Section 2, we study the exponential family $(F_\theta)_{\theta\geq 0}$. We give a heuristic proof that F_θ can be approximated by a lognormal distribution and obtain an approximation of the Legendre-Fenchel transform of the lognormal distribution. The first important application of our results, namely the saddlepoint approximation for $\mathbb{P}(S_n \leq x)$, is given in Section 3. The second is a Monte Carlo estimator for $\mathbb{P}(S_n \leq x)$ given in Section 4.2. It follows a classical model (Asmussen & Glynn [6, VI.2]) by attempting importance sampling with importance distribution is $F_{\theta(x)}$, but the implementation faces the difficulty that neither $\theta(x)$ nor $\kappa\big(\theta(x)\big)$ are explicit, but must be approximated (various approaches to overcome this are discussed in the numerical examples presented in Section 5). The algorithm requires simulation from F_θ for certain θ, and we suggest an acceptance-rejection (A-R) for this with a certain Gamma proposal; the analysis gives as byproduct that this Gamma is an excellent approximation of F_θ. The Appendix contains various supplements, in particular a proof that the importance sampling proposed in Section 4.2 has a certain asymptotical efficiency property.

2 The exponential family generated by the lognormal distribution

We let F be the cumulative distribution function of X and adopt the notation $X \sim \mathrm{LN}(0, \sigma^2)$. For convenience, we write f_n and F_n for the pdf and cdf of S_n, respectively.

The exponential tilting scheme in the Introduction is often also referred to as *Esscher transformation*. Note that since $\kappa(\theta)$ is well-defined for all $\theta > 0$, one avoids for $x < \mathbb{E}X$ (the relevant case for our left tail problem) the difficulties in large deviations theory associated with boundary problems when defining the saddlepoint $\theta(x)$ and which lead into minimizing the convex conjugate $\kappa^*(x) = \kappa(\theta(x)) + x\theta(x)$ (also called the *Legendre-Fenchel transform*).

Proposition 2.1. *Let $X \sim \mathrm{LN}(0, \sigma^2)$, $k \in \mathbb{N}^+$, $\theta > 0$. Then*

$$\mathbb{E}[X^k \mathrm{e}^{-\theta X}] = \mathbb{E}\big[g_k(\sigma_{k,\theta} Z)\big] \frac{\exp\left\{ - \dfrac{\mathcal{W}^2(\theta\sigma^2 \mathrm{e}^{k\sigma^2}) + 2\,\mathcal{W}(\theta\sigma^2 \mathrm{e}^{k\sigma^2}) - k^2\sigma^4}{2\sigma^2} \right\}}{\sqrt{1 + \mathcal{W}(\theta\sigma^2 \mathrm{e}^{k\sigma^2})}}. \quad (2.1)$$

In addition, also

$$\mathbb{E}[X^k \mathrm{e}^{-\theta X}] = \mathbb{E}\big[g_0\big((Z+k)\sigma_{0,\theta}^2\big)\big] \times \frac{\exp\left\{ - \dfrac{\mathcal{W}^2(\theta\sigma^2) + 2\,\mathcal{W}(\theta\sigma^2)}{2\sigma^2} + k\big(\mathcal{W}(\theta\sigma^2) + k\sigma_{0,\theta}^2/2\big) \right\}}{\sqrt{\mathcal{W}(\theta\sigma^2) + 1}}, \quad (2.2)$$

where $Z \sim \mathrm{N}(0,1)$ and

$$\sigma_{k,\theta}^2 = \frac{\sigma^2}{1 + \mathcal{W}(\theta\sigma^2 \mathrm{e}^{k\sigma^2})}, \quad g_k(w) = \exp\left\{ - \frac{\mathcal{W}(\theta\sigma^2 \mathrm{e}^{k\sigma^2})}{\sigma^2}\big(\mathrm{e}^w - 1 - w - w^2/2\big) \right\}.$$

Here

$$\mathbb{E}[g_k(\sigma_{k,\theta} Z)] = 1 + O\big((\log\theta)^{-1}\big), \qquad \mathbb{E}[g_0(\sigma_{k,\theta} Z)] = 1 + O\big((\log\theta)^{-2}\big). \quad (2.3)$$

The proof of Proposition 2.1 follows the same lines as the proof of Proposition 2.1 in the companion paper [7] and therefore omitted. We just note here that the approximations (2.1) and (2.2) are obtained by applying the Laplace method (cf. [19]). Roughly speaking the Laplace method employs a second order expansion of the exponent of the integrand defining $\mathbb{E}[X^k \mathrm{e}^{-\theta X}]$ around a specific value. The first approximation is obtained using the standard Laplace's method with an expansion around the value maximizing the exponent $\rho_k = \mathcal{W}(\theta\sigma^2 \mathrm{e}^{k\sigma^2})$ while the second one is obtained instead by using an expansion around the value $\rho = -\mathcal{W}(\theta\sigma^2)$.

Next we focus on finding asymptotic approximations for the derivatives of the cumulant transform $\kappa(\theta)$. Recall that such derivatives are associated to the Esscher transform and the distributions in the exponential family generated by the lognormal via (1.1). Combining this with Proposition 2.1, we arrive at the following asymptotic equivalences:

Corollary 2.2. *Let $X \sim \mathrm{LN}(0, \sigma^2)$ and define*

$$\mu_\theta = -\mathcal{W}(\theta\sigma^2), \qquad \sigma_\theta^2 = \frac{\sigma^2}{1 + \mathcal{W}(\theta\sigma^2)}. \tag{2.4}$$

Then

$$\lim_{\theta\to\infty} \frac{\mathbb{E}_\theta[X]}{\exp\{\mu_\theta + \sigma_\theta^2/2\}} = 1, \qquad \lim_{\theta\to\infty} \frac{\mathrm{Var}_\theta[X]}{\exp\{2\mu_\theta + \sigma_\theta^2\}(\mathrm{e}^{\sigma_\theta^2} - 1)} = 1.$$

where $\mathbb{E}_\theta$ and Var_θ are the expectation and variance operators under F_θ.

Proof. Using (2.2) in Proposition 2.1 we arrive at

$$\mathbb{E}_\theta[X] = \exp\{\mu_\theta + \sigma_\theta^2/2\} \frac{\mathbb{E}\big[g_1((Z+1)\sigma_{1,\theta})\big]}{\mathbb{E}\big[g_0(Z\sigma_{0,\theta})\big]}$$

and

$$\mathrm{Var}_\theta[X] = \exp\{2\mu_\theta + \sigma_\theta^2\} \left[\mathrm{e}^{\sigma_\theta^2} \frac{\mathbb{E}\big[g_2((Z+2)\sigma_{2,\theta})\big]}{\mathbb{E}\big[g_0(Z\sigma_{0,\theta})\big]} - \frac{\mathbb{E}^2\big[g_1((Z+1)\sigma_{1,\theta})\big]}{\mathbb{E}^2\big[g_0(Z\sigma_{0,\theta})\big]}\right]$$

where $Z \sim \mathrm{N}(0,1)$ and the function g_0, g_1, g_2 are defined as in Proposition 2.1. The result follows as a consequence of (2.3). □

Interestingly, one could at least at the heuristic level identify F_θ as approximate lognormal with parameters

$$\mu_\theta = -\mathcal{W}(\theta\sigma^2), \qquad \sigma_\theta^2 = \frac{\sigma^2}{1 + \mathcal{W}(\theta\sigma^2)}.$$

Thus we conclude that a sensible approximation of the Esscher transform of a lognormal distribution is again a lognormal distribution whose parameters are given explicitly in terms of the Lambert W function. Moreover, the expectation and variance of such a lognormal random variable coincide with the values given in formula (2.4).

The argument is to use the Laplace method to get

$$F_\theta(x) = \frac{\mathbb{E}[\mathrm{e}^{-\theta X}\,\mathbb{I}(X < x)]}{\mathbb{E}[\mathrm{e}^{-\theta X}]} = \frac{\mathbb{E}\big[g_0(W)\,\mathbb{I}(W < \log x - \mu_\theta)\big)\big]}{\mathbb{E}[g_0(W)]}, \qquad x > 0,$$

where $W \sim \mathrm{N}(0, \sigma_\theta^2)$ and g_0 is defined as in Proposition 2.1. The Laplace method constructs the approximation above in such way that the function g_0 is close enough to 1 around a neighborhood of 0 where the mass of the random variable W is concentrated. Hence, neglecting the error associated to the function g_0 and normalizing we arrive at

$$F_\theta(x) \approx \mathbb{P}\big(W < \log x - \mu_\theta\big).$$

In Fig. 1, the right solid line blue plot is the $\mathrm{LN}(0, \sigma^2)$ density with $\sigma = 0.25$. The other solid line blue plots are (from right to left) the F_θ-densities for $\theta = 10, 25, 100$. It is notable how little even such a large values as $\theta = 100$ shifts the distribution towards the origin, which can be explained by the lognormal density decaying only slowly to 0 as $x \downarrow 0$. The two dotted red plots are the lognormal approximations of the F_θ-densities for $\theta = 25$ and 100. In Section 4.1 we derive an alternative approximation in terms of the Gamma distribution.

Figure 1: F_θ-densities and their lognormal approximations

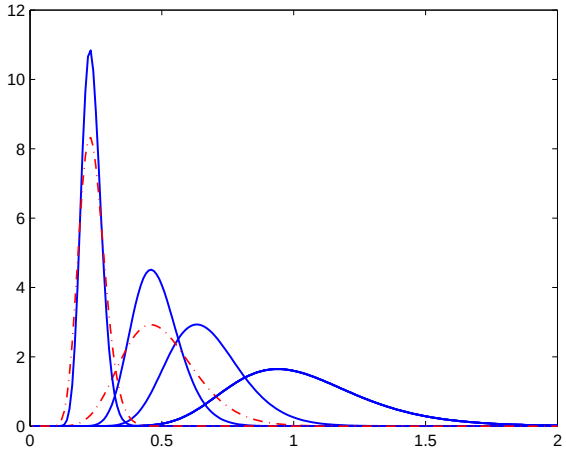

2.1 The Cramér function

The previous results will allow us to provide asymptotic approximations for the Cramér function of the lognormal distribution $\theta(x)$, i.e. the solution of the equation $\mathbb{E}_{\theta(x)}[X] = x$. Our suggestion is

$$\widetilde{\theta}(x) = \frac{\gamma(x)\mathrm{e}^{\gamma(x)}}{\sigma^2} \qquad \text{where} \qquad \gamma(x) = \frac{-1 - \log x + \sqrt{(1-\log x)^2 + 2\sigma^2}}{2}. \tag{2.5}$$

To arrive at (2.5), we approximate the value $\theta(x)$ solving the equation $\mathbb{E}_\theta[X] = x$ with the solution $\widetilde{\theta}(x)$ of

$$\exp\{\mu_\theta + \sigma_\theta^2/2\} = x \tag{2.6}$$

with μ_θ and σ_θ^2 as given in (2.4). Writing $\mathcal{W} = \mathcal{W}(\theta\sigma^2)$ for the moment, (2.6) then means

$$-\mathcal{W} + \frac{\sigma^2/2}{1+\mathcal{W}} \;=\; \log x$$

or equivalently that $\gamma = \mathcal{W}$ is the solution of the quadratic $\gamma^2 + (1+\log x)\gamma - \sigma^2/2 + \log x = 0$. This gives (2.5) (excluding the negative sign of the square root by an easy argument).

By easy calculus,

$$\gamma(x) = -\log x + o(1) \;\text{ as } x \to 0. \tag{2.7}$$

In consequence, the approximated solution $\widetilde{\theta}(x)$ is asymptotically equivalent to

$$\widetilde{\theta}(x) \sim -\frac{\log x}{x\sigma^2}, \qquad x \to 0.$$

In Sections 3 and 4.2 we will employ these results to construct a saddlepoint approximation and a Monte Carlo estimator of the left tail probability of a sum of lognormal random variables. In particular, the asymptotic results derived above will be useful to show that the approximation $\widetilde{\theta}(x)$ is asymptotically sharp and such that when used as the twisting parameter of an exponential change of measure estimator it remains asymptotically efficient as $x \to 0$.

3 Saddlepoint approximation in the left tail of a lognormal sum

Daniels' saddlepoint method produces an approximation of the density function of a sum of i.i.d. random variables which is valid asymptotically on the number of summands. The first and second order approximations are embodied in the formula

$$f_n(nx) \approx \sqrt{\frac{1}{2\pi n\kappa''(\theta(x))}}\exp\big\{n\kappa^*(x)\big\}\big(1+\frac{1}{n}[\zeta_4(\theta(x))/8+5\zeta_3(\theta(x))^2/24]\big),$$

where $\kappa^*(x)=\kappa(\theta(x))+x\theta(x)$ is the convex conjugate of $\kappa(x)$ and

$$\zeta_k(\theta)=\frac{\kappa^{(k)}(\theta)}{\kappa''(\theta)^{k/2}},$$

is the standardized cumulant.

The corresponding saddlepoint approximation for the cumulative distribution function is given by (Jensen, 1995, [19])

$$\begin{aligned}F_n(nx) = \frac{1}{\lambda_n(x)}\exp\big\{n\kappa^*(x)\big\}\Big\{&B_0(\lambda_n(x))\\ &+\frac{\zeta_3(\theta(x))}{6\sqrt{n}}B_3(\lambda_n(x))+\frac{\zeta_4(\theta(x))}{24n}B_4(\lambda_n(x))+\frac{\zeta_3(\theta(x))^2}{72n}B_6(\lambda_n(x))\Big\},\end{aligned}$$

where

$$\lambda_n(x)=\theta(x)\sqrt{n\kappa''(\theta(x))},\quad B_0(\lambda)=\lambda \mathrm{e}^{\lambda^2/2}\Phi(-\lambda),$$

and

$$\begin{aligned}B_3(\lambda)&=-\big\{\lambda^3B_0(\lambda)-(\lambda^3-\lambda)/\sqrt{2\pi}\big\},\\ B_4(\lambda)&=\lambda^4B_0(\lambda)-(\lambda^4-\lambda^2)/\sqrt{2\pi},\\ B_6(\lambda)&=\lambda^6B_0(\lambda)-(\lambda^6-\lambda^4+3\lambda^2)/\sqrt{2\pi}.\end{aligned}$$

General results for the saddlepoint approximation state that for a fixed x the relative error is $O(1/n)$ for the first order approximation and $O(1/n^2)$ for the second order approximation. More can be said, however, for the case of a lognormal sum. It is simple to see that the density $f(x)$ is logconcave for $x<e^{1-\sigma^2}$ (second derivative of $\log(f(x))$ is negative) and according to Jensen (1995, section 6.2) [19] we therefore have that the saddlepoint approximations have the stated relative errors uniformly for x in a region around zero. Furthermore, write the exponentially tilted density as $\exp(-h(x)-\kappa(\theta))$ with $h(x)=\log(x)+(\log(x))^2/(2\sigma^2)+\theta x$. We center and scale this density using $w=\mathcal{W}(\theta\sigma^2)$ as follows

$$\begin{aligned}h_0(u)&=h(e^{-w}(1+\frac{\sigma}{\sqrt{w}}u))-h(e^{-w})\\ &=\log(1+\frac{\sigma}{\sqrt{w}}u)+\frac{1}{2\sigma^2}\{[-w+\log(1+\frac{\sigma}{\sqrt{w}}u)]^2-w^2\}+\frac{\sqrt{w}}{\sigma}u.\end{aligned}$$

From this we find $h_0(u) = \frac{1}{2}u^2 + O(\sigma(|u| + |u|^3)/\sqrt{w})$. Thus, for $w \to \infty$ ($\theta \to \infty$), the density converges uniformly in the region $|u| < (\sqrt{w}/\sigma)^{1/6}$, say, to the standard normal density. Due to the logconcavity the left tail beyond $-(\sqrt{w}/\sigma)^{1/6}$ is well behaved and for the right tail we find $h_0'(u) > (\sqrt{w}/\sigma)^{1/6}/2$ for $\sigma/\sqrt{w} < 1$. The convergence to the standard normal density as $w \to \infty$ implies that the saddlepoint approximations become exact in the limit $w \to \infty$.

To evaluate the saddlepoint approximation we need to calculate the Laplace transform and its derivatives numerically. We want to implement the integration in such a way that the relative accuracy of the integration is of the same order irrespective of the argument θ. For $k = 0, 1, 2, 3, 4$ we want to evaluate the integral

$$L_k(\theta) = E(X^k \exp(-\theta X)) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-h(y))dy, \quad h(y) = \theta e^y + \frac{1}{2\sigma^2}y^2 - ky$$

The minimum of $h(y)$ is attained at $y_0 = k\sigma^2 - w$, where $w = \mathcal{W}(\theta\sigma^2 \exp(k\sigma^2)) > 0$ and $h(y_0) = w^2/(2\sigma^2) + w/\sigma^2 - \frac{1}{2}\sigma^2 k^2$. Furthermore,

$$h_0(z) = h(y_0 + z) - h(y_0) = \frac{w}{\sigma^2}(e^z - 1 - z) + \frac{z^2}{2\sigma^2}.$$

Since $h_0''(z) = (we^z + 1)/\sigma^2 > 0$ we see that h is convex. Choosing a scale τ such that $2h_0(-\tau) = 1$ we obtain that $h_1(u) = 2h_0(\tau u)$ is a convex function bounded between 0 and 1 for $-1 < u < 0$, is above $-u$ for $u < -1$ and with $h_1(u) \geq h_1(-u)$ for $u > 0$. In this way the precision of the numerical integration of $\exp(-\frac{1}{2}h_1(u))$ will be of the same order irrespective of the value of w and σ^2. In practice we can take τ as

$$\tau = \begin{cases} \dfrac{\sigma}{\sqrt{1+w}}, & \dfrac{\sigma}{\sqrt{1+w}} \leq c_0, \\ \sqrt{w^2 + 2w + \sigma^2} - w, & \dfrac{\sigma}{\sqrt{1+w}} > c_0, \end{cases}$$

where c_0 is an arbitrary constant. Unless σ^2 is large we can use $\tau = \sigma/\sqrt{1+w}$ for all w.

The saddlepoint $\theta(x)$ being the solution to $\kappa'(\theta(x)) = -x$ can be found by Newton-Raphson iteration using $\tilde{\theta}(x)$ as the initial value. In Table 1 the initial value $\tilde{\theta}(x)$ is given together with $\theta(x)$ and the mean value $E_{\tilde{\theta}(x)}(X)$. In all cases in the table four steps in the Newton-Raphson suffices for reaching $\theta(x)$.

In Tables 2 and 3 are examples with the saddlepoint approximation to the left tail probability.

Table 1: Comparison of initial value $\tilde{\theta}(x)$ and final value $\theta(x)$ for the case $\sigma = 0.250$.

x	$\mathbb{E}_{\widetilde{\theta}(x)}[X]$	$\tilde{\theta}(x)$	$\theta(x)$
1.0	0.99905160	0.5002255	0.4850103
0.9	0.89695877	2.4295388	2.3625893
0.8	0.79589537	5.0894397	4.9624633
0.7	0.69554784	8.8690980	8.6691868
0.5	0.49617443	23.1845282	22.7639315
0.3	0.29767635	65.8850274	64.9626105
0.1	0.09934273	373.4301331	369.9235664

Table 2: Approximation of the CDF of a lognormal sum with $n = 4$ and $\sigma = 0.25$.

x	$\theta(x)$	Saddle0	Saddle1	MC/AA	MC/MC
0.6500	11.132319	0.0001536084	0.0001592339	1.61e-04 ± 1.76e-06	1.61e-04 ± 1.83e-06
0.7000	8.669187	0.0012499087	0.0013015022	1.31e-03 ± 1.34e-05	1.30e-03 ± 1.37e-05
0.7500	6.644334	0.0065782847	0.0068830734	6.93e-03 ± 6.50e-05	6.92e-03 ± 6.66e-05
0.8000	4.962463	0.0242679549	0.0255206432	2.56e-02 ± 2.20e-04	2.55e-02 ± 2.24e-04
0.8500	3.552969	0.0669477011	0.0707464921	7.13e-02 ± 5.53e-04	7.11e-02 ± 5.61e-04
0.9000	2.362589	0.1456850237	0.1545557418	1.55e-01 ± 1.09e-03	1.55e-01 ± 1.10e-03

Table 3: Approximation of the CDF of a lognormal sum with $n = 64$ and $\sigma = 0.25$.

x	$\theta(x)$	Saddle0	Saddle1
0.90	2.3625893	8.693420e-06	8.772302e-06
0.91	2.1470381	3.951385e-05	3.989503e-05
0.92	1.9383125	1.575592e-04	1.591772e-04
0.93	1.7361482	5.538798e-04	5.599406e-04
0.95	1.3505093	4.782814e-03	4.842303e-03
0.97	0.9882486	2.646345e-02	2.683567e-02
0.99	0.6476640	9.774927e-02	9.926919e-02

4 Simulation

4.1 Random variate generation from F_θ

We first consider the problem of generating a r.v. from the density

$$f_\theta(x) = \mathrm{e}^{-\theta x - \kappa(\theta)} f(x), \qquad x > 0.$$

The obvious naive choice is acceptance-rejection (A-R; [6, II.2]), simulating Z from f and rejecting w.p. $\mathrm{e}^{-\theta Z}$. This choice produces a very simple algorithm for generating from f_θ and the method is exact even when we do not have an explicit expression for $\kappa(\theta)$.

Algorithm 4.1. Simulate $U \sim \mathrm{U}(0,1)$, $Z \sim \mathrm{LN}(0, \sigma^2)$. If $U > \mathrm{e}^{-\theta Z}$ repeat. Else, return Z.

Ideally, we would like to have a rejection probability p as close to 1 but in our case $p = \mathrm{e}^{\kappa(\theta)}$, so as the value of θ increases, the probability of acceptance diminishes and hence the expected number c of rejection steps goes to ∞. In consequence, this estimator is very inefficient for large values of θ.

An alternative algorithm is as follows. Observe that $f_\theta(x)$ is proportional to $\exp\{-\theta x - \log^2 x)/2\sigma^2 - \log x\}$ which is maximized at $x = m = w/\theta\sigma^2$ where $w = \mathcal{W}(\theta\sigma^2 \mathrm{e}^{-\sigma^2})$. Let $Y_\theta \sim f_\cdot$ Up to a constant, the density of Y_θ/m is

$$\begin{aligned} f(mx) &= \exp\{-\theta m x - \log^2(mx)/2\sigma^2 - \log(mx)\} \\ &\propto \exp\{-\sigma^{-2}[wx + \log^2 m/2 + \log^2(x)/2 + \log m \log x + \sigma^2 \log x]\} \\ &\propto \exp\{-\sigma^{-2}[wx + \log^2(x)/2 - w \log x]\} \\ &= x^{\sigma^{-2} w} \mathrm{e}^{-\sigma^{-2} w x} \mathrm{e}^{-\log^2(x)/2\sigma^2}. \end{aligned} \tag{4.1}$$

Note that $-w = \log m + \sigma^2$. This gives the following A-R algorithm:

Algorithm 4.2. Simulate $U \sim \mathrm{U}(0,1)$ and $Z \sim \mathrm{Gamma}(w/\sigma^2 + 1, w/\sigma^2)$. If $U > \mathrm{e}^{-\log^2(Z)/2\sigma^2}$ repeat. Else, return $Y_\theta = mZ = wZ/\theta\sigma^2$.

Now the mean of Z is $1 + \sigma^2/w$ and the variance is $\omega_\theta^2 = (1 + \sigma^2/w)\sigma^2/w$. Since $\omega_\theta^2 \to 0$ as $\theta \to \infty$ and hence $w \to \infty$, $Z \to 1$ in probability. Hence, by dominated convergence, the overall probability $\mathbb{E}[\mathrm{e}^{-\log^2(Z)/2\sigma^2}]$ converges to 1!

Remark 4.3. (Gamma approximation) *A-R* with acceptance probability $\to 1$ is only possible if the ratio of the proposal and target density goes to 1. Thus, Y_θ is accurately approximated by $wZ/\theta\sigma^2 \sim \mathrm{Gamma}(w/\sigma^2 + 1, \theta)$. And if we pass to sums, $S_n = Y_{\theta,1} + \cdots + Y_{\theta,n}$ is accurately approximated by the Gamma r.v. $w(Z_1 + \cdots + Z_n)/\theta\sigma^2$ which is $\mathrm{Gamma}(n(w/\sigma^2 + 1), \theta)$.

Remark 4.4. (Normal approximation) Note that to first order, $w \sim \log\theta$. Thus the variance $(w+\sigma^2)/\theta^2\sigma^2$ goes to 0, and we can approximate Y by a r.v. distributed as $\mathrm{N}\big((w+\sigma^2)/\theta\sigma^2), (w+\sigma^2)/\theta^2\sigma^2\big)$ and S_n with $\mathrm{N}\big(n(w+\sigma^2)/\theta\sigma^2), n(w+\sigma^2)/\theta^2\sigma^2\big)$.

When plotted on top of Fig. 1, the two Gamma approximations for $\theta = 25$ and 100 are indistinguishable from the lognormal$(\mu_\theta, \sigma_\theta^2)$ approximations.

4.2 Efficient Monte Carlo for left tails of lognormal sums

In this section we develop an *asymptotically efficient Monte Carlo estimator* $\widehat{\alpha}_n(x)$, for the left tail probability of a lognormal sum $\alpha_n(x) = \mathbb{P}(S_n \leq nx)$ as $x \to 0$.

We start by recalling some standard concepts from rare event simulation ([6, VI.1]). In our setting, we say that a Monte Carlo estimator $\widehat{\alpha}_n(x)$ is *strongly efficient* or has *bounded relative error* as $u \to 0$ if

$$\limsup_{x \to 0} \frac{\operatorname{Var} \widehat{\alpha}_n(x)}{\alpha_n^2(x)} < \infty.$$

This efficiency property implies that the number of replications required to estimate $\alpha_n(nx)$ with certain fixed relative precision remains bounded as $x \to 0$. A weaker criterion is *logarithmic efficiency* defined as

$$\limsup_{x \to 0} \frac{\operatorname{Var} \widehat{\alpha}_n(x)}{\alpha_n^{2-\epsilon}(x)} = 0, \qquad \forall \epsilon > 0.$$

From a practical point of view, there is no substantial difference between these two criteria. However, it is often easier to prove logarithmic efficiency rather than bounded relative error. Logarithmic efficiency implies that the number of replications needed for achieving certain relative precision grows at rate of order at most $|\log(\alpha_n(x))|$.

An alternative unbiased estimator can be obtained by using the variance reduction technique *importance sampling* ([6, V.1]). This method relies on the existence of a Radon-Nikodym derivative with respect to a probability measure, say $\mathbb{Q}$. If we are interested in estimating $\mathbb{E}[h(W)]$ where W is random, h an measurable function, $\mathbb{E}$ is the expectation operator under the measure $\mathbb{P}$ and $\mathbb{Q}$ is an absolutely continuous measure with respect to $\mathbb{P}$, then it holds that

$$\mathbb{E}[h(W)] = \mathbb{E}^{\mathbb{Q}}[L\, h(W)],$$

where $\mathbb{E}^{\mathbb{Q}}$ is the expectation operator under the measure $\mathbb{Q}$ and $L = \mathrm{d}\mathbb{P}/\mathrm{d}\mathbb{Q}$ is the Radon-Nikodym derivative of $\mathbb{P}$ with respect to $\mathbb{Q}$ (the last also goes under the name *likelihood ratio* in the simulation community). Hence, if X is simulated according to $\mathbb{Q}$, then $L\, h(W)$ serves as an unbiased estimator of the quantity $\mathbb{E}[h(W)]$. The strategy of selecting an importance distribution from the exponential family generated by the lognormal $\{F_\theta : \theta \in \Theta\}$ is often referred as *exponential twisting*, *exponential tilting* or simply *exponential change of measure*. Ideally, the twisting parameter θ is selected as the value of the Cramér function $\theta(\cdot)$ evaluated at x, and defined via

$$-\kappa'(\theta(x)) = \mathbb{E}_{\theta(u)}[X] = x. \tag{4.2}$$

Doing so, exponential twisting is logarithmically efficient as $n \to \infty$, cf. [6, p. 169-171].

Notice, however, that difficulties arise in the right tail if the X_i's are heavy-tailed: then the integral associated with $\mathbb{E}[\mathrm{e}^{-\theta X}]$ diverges for negative values of the argument θ and in consequence, the equation (4.2) has no solution if $x > \mathbb{E}[X_i]$. Further difficulties in the heavy-tailed environment are exposed in [4, 9]. Nevertheless, exponential twisting can be implemented for the left tail probability of a lognormal sum; moreover, it turns out that it is logarithmically efficient:

Theorem 4.5. *Consider $X_1, \ldots, X_n \sim F_{\theta(u)}$ and set $S_n = X_1 + \cdots + X_n$. Define*

$$\beta_n(x) := \mathcal{L}^n(\theta(x))\, \mathrm{e}^{\theta(x) S_n}\, \mathbb{I}\{S_n < nx\},$$

where $\mathcal{L}(\cdot)$ is the Laplace transform of the lognormal distribution. Then $\beta_n(x)$ is a logarithmically efficient and unbiased estimator of $\alpha_n(x)$ as $n \to \infty$.

Proof. The lognormal density is log-concave and so the result follows immediately from the proof of Theorem 2.10, Chapter VI in [6]. □

Notice however, that the optimal exponential twisting algorithm described above is not implementable since the Laplace transform $\mathcal{L}(\theta)$ and the Cramér function $\theta(\cdot)$ are unknown. In this paper, we propose an alternative estimator of $\alpha_n(x)$ which is logarithmic efficient as $x \to 0$ (not necessarily as $n \to \infty$!). For its construction we employ the approximation of the Cramér function provided in the previous section and we assume that an unbiased estimator of the Laplace transform is available. The algorithm is as follows:

Algorithm 4.6.

1. Use the approximation $\widetilde{\theta}(x)$ of the Cramér function given in (2.5).
2. Obtain an unbiased estimate $\widehat{\mathcal{L}}(\widetilde{\theta}(x))$ of the Laplace transform [cf. 7]
3. Simulate $X_1, \ldots, X_n \sim F_{\widetilde{\theta}(x)}$ and set $S_n = X_1 + \cdots + X_n$.
4. Return
$$\widehat{\alpha}_n(x) = \mathrm{e}^{\widetilde{\theta}(x) S_n} \left[\widehat{\mathcal{L}}(\widetilde{\theta}(x))\right]^n \mathbb{I}(S_n < nx). \tag{4.3}$$

However, even if $\widehat{\mathcal{L}}(\theta)$ was unbiased, then the estimator (4.3) is biased for $\alpha_n(x)$; because of Jensen's inequality it holds that $\mathbb{E}[\widehat{\mathcal{L}}^n(\theta)] > \mathcal{L}^n(\theta)$. A solution is to take the product of n independent copies of $\widehat{\mathcal{L}}(\theta)$ to estimate $\mathcal{L}^n(\theta)$ without bias. We go more into this topic in Section 5 and Section B of the Appendix.

We discuss next an interesting asymptotic efficiency property of (4.3) when we employ the logarithmically efficient unbiased estimator of $\mathcal{L}(\theta)$ suggested in [7]. That is

$$\widehat{\mathcal{L}}_{IS}(\theta) = \widetilde{\mathcal{L}}(\theta)\, \vartheta(Y, \theta), \tag{4.4}$$

where $Y \sim \mathrm{N}(0, \sigma^2)$,

$$\widetilde{\mathcal{L}}(\theta) = \exp\left\{-\frac{\mathcal{W}^2(\theta\sigma^2) + 2\,\mathcal{W}(\theta\sigma^2)}{2\sigma^2}\right\}, \qquad \vartheta(t;\theta) = \exp\left\{-\frac{\mathcal{W}(\theta\sigma^2)}{\sigma^2}\big(\mathrm{e}^t - 1 - t\big)\right\}.$$

Proposition 4.7. *Let*

$$\widehat{\alpha}_n(x) = \mathrm{e}^{\widetilde{\theta}(x) S_n} \left[\widehat{\mathcal{L}}_{\mathrm{IS}}(\widetilde{\theta}(x)\right]^n \mathbb{I}(S_n < nx) \tag{4.5}$$

Then $\widehat{\alpha}_n(x)$ is a consistent estimator of $\alpha_n(x)$. Moreover, it is also a logarithmic efficient estimator in the mean square error sense. That is, for all $\epsilon > 0$

$$\limsup_{x \to 0} \frac{\mathrm{MSE}_{\alpha_n(x)}[\widehat{\alpha}_n(x)]}{\alpha_n^{2-\epsilon}(x)} = 0, \tag{4.6}$$

For the proof, see Section C of the Appendix.

4.3 Density estimation

Consider the problem of estimating the density of a lognormal sum via simulation. Following [6], Example V.4.3 p. 146, slightly extended, we first note that the conditional density at nx of S_n given

$$S_{n,-i} \;=\; X_1 + \cdots + X_{i-1} + X_{i+1} + \cdots + X_n \;=\; S_n - X_i$$

is $f(nx - S_{n,-i})$. Hence an unbiased estimator of $f_n(nx)$ is $\sum_1^n f(nx - S_{n,-i})/n$.

However, since we are dealing with values of x far to the left of $\mathbb{E}X$, it is likely that $S_{n,-i} > x$ so that $f(nx - S_{n,-i}) = 0$ and the procedure will come out with a large number of zeroes. Hence we employ the same importance sampling estimator as used elsewhere. That is, we simulate the X_j from $f_{\tilde{\theta}(x)}$ and return the estimator

$$\widehat{f}_n(nx) \;=\; \frac{\exp\{\tilde{\theta}(x)S_n + n\kappa\big(\tilde{\theta}(x)\big)\}}{n} \sum_{i=1}^{n} f(nx - S_{n,-i})$$

(in practice to be averaged over R replications). An alternative slightly more complicated estimator is

$$\widehat{f}_n(nx) \;=\; \frac{1}{n} \sum_{i=1}^{n} f(nx - S_{n,-i}) \exp\{\tilde{\theta}(x)S_{n,-i} + (n-1)\kappa\big(\tilde{\theta}(x)\big)\}\,.$$

In [18], an importance sampling estimator for $F_n(z)$ is suggested and it is written that a parallel estimator for $f_n(z)$ can be constructed in the same way. We do not follow the details of this statement.

5 Numerical examples

In our numerical experiments, we have taken parameter values that we consider realistic from the point of view of financial applications. A yearly volatility of order 0.25 is often argued to be typical. We have considered periods of lengths one year, one quarter, one month and one week, corresponding to $\sigma = 0.25$, $\sigma = 0.25/\sqrt{4} = 0.125$, $\sigma = 0.25/\sqrt{12} = 0.072$, resp. $\sigma = 0.25/\sqrt{52} = 0.035$. Real-life portfolios are often large, even in the thousands; the values we have chosen are $n = 4, 16, 64, 256$.

For each combination of n and σ we have conducted several numerical empirical analyses. In all numerical experiments involving simulation we have employed $R = 100,000$ replications.

Transformations associated to the Laplace transform

We test empirically the approximations for the n-th power of the Laplace transform $\mathcal{L}^n(\theta)$. The approximations discussed here are alternatives to numerical integration. We considered two approximations: the n-th power of the approximation derived from (1.2); and the n-th power of the IS sampling estimator (4.4). Notice however that the last estimator is biased for any $n > 1$ so the bias will grow exponentially with n. To address this issue we considered two alternatives (see Appendix B for

Table 4: Approximated values of the n-th power of the Laplace transform with $n = 256$ and $\sigma = 0.250$.

θ	$\widetilde{\mathcal{L}^n}(\theta)$	$\widehat{\mathcal{L}^n}(\theta)$	$\widehat{\mathcal{L}^n}(\theta)$ – Correction	$\widehat{\mathcal{L}^n}(\theta)$
0.9705	1.23e-108	1.10e-108±3.14e-111	1.10e-108±3.14e-111	1.12e-108±4.92e-111
0.9010	4.10e-101	3.82e-101±9.86e-104	3.82e-101±9.86e-104	3.75e-101±1.53e-103
0.8322	1.20e-93	1.08e-93±2.71e-96	1.08e-93±2.71e-96	1.10e-93±4.11e-96
0.7642	3.08e-86	2.84e-86±6.49e-89	2.84e-86±6.49e-89	2.87e-86±9.84e-89
0.6971	6.95e-79	6.85e-79±1.35e-81	6.85e-79±1.35e-81	6.49e-79±2.02e-81
0.6307	1.38e-71	1.27e-71±2.45e-74	1.27e-71±2.45e-74	1.29e-71±3.63e-74
0.5651	2.40e-64	2.31e-64±3.89e-67	2.31e-64±3.89e-67	2.26e-64±5.71e-67
0.5002	3.69e-57	3.45e-57±5.36e-60	3.45e-57±5.36e-60	3.49e-57±7.83e-60

further details): insert a bias correction term and to consider the unbiased estimator built as the product of n independent copies of (4.4).

We consider an example with large $n = 256$ and $\sigma = 0.25$; the number of replications for all estimators was $R = 100,000$. The results can be found in Table 4. It is notorious that for moderate values of the parameter θ one obtains very small values of the Laplace transform. Also, the numerical results of this example indicate that the approximation $\widetilde{L}(\theta)$ underestimates the real value of $\mathcal{L}(\theta)$. It is also noted that the Bias Correction (BC) does not provide a significant improvement over the value of the approximation. In particular, the variance of the estimator $\widehat{\mathcal{L}^n}(\cdot)$ is only slightly larger that the variance of $\widehat{\mathcal{L}}^n(\cdot)$ and its bias appears to be very small. In spite of these minor pitfalls, we consider all the estimators to be very sharp. However, $\widehat{\mathcal{L}^n}(\cdot)$ is unbiased so the undesired amplifying effect on the bias produced by the n-th power transformation is avoided. We favor the use of the estimators $\widehat{\mathcal{L}^n}(\cdot)$ because of its unbiasedness and relatively small variance.

Left tail of the Lognormal Sum

Next we verify the approximations for the cdf and pdf of the lognormal sum. We have thereby been thinking of a portfolio of n assets with next-period values $Y_1, \ldots, Y_n$ assumed i.i.d. lognormal(μ, σ^2), such that a loss corresponds to a small value x of $S_n = Y_1, \ldots, Y_n$. When choosing x, we have had the recommended VaR values 0.99%–0.99.97% of Basel II [1] in mind and chosen $\mathbb{P}(S_n \leq nx)$ to be in the interval 0.0001–0.0100.

We have proposed two type of approximations: saddlepoint approximations and Monte Carlo estimators. We start with the saddlepoint approximation for both the pdf and cdf of the sum of n lognormals and which are given by

$$\widetilde{F_n}(nx) = \exp\left\{n\kappa^*(x)\right\} \mathrm{e}^{\lambda_n(x)^2/2}\, \Phi(-\lambda_n(x)), \qquad \widetilde{f}_n(nx) = \frac{\exp\left\{n\kappa^*(x)\right\}}{\sqrt{2\pi n\kappa''(\theta(x))}},$$

where $\lambda_n(x) = \theta(x)\sqrt{n\kappa''(\theta(x))}$ and $\kappa^*(\cdot)$ is the complex conjugate of $\kappa(\cdot)$. Recall that $\exp\{n\kappa^*(x)\} = \mathcal{L}^n(\theta(x))\mathrm{e}^{x\theta(x)}$. In our numerical results we report saddlepoint approximations (labeled Saddle/MC); and 2) Monte Carlo estimators (labeled MC/MC). The last is based on the proposed importance sampling estimator where

the importance sampling distribution was selected to be from the exponential family. The parameter θ defining such distribution was selected to be equal to the approximation of the Cramér function $\widetilde{\theta}(\cdot)$ evaluated at x/n. The general estimator for the CDF of the lognormal sum has the form

$$\widehat{F}_n(nx) = \mathcal{L}^n(\widetilde{\theta}(x))\, \mathrm{e}^{\widetilde{\theta}(x)S_n}\, \mathbb{I}\{S_n < nx\},$$

where $S_n = X_1 + \cdots + X_n$ and $X_1, \ldots, X_n$ is a sample from the exponential family. Similarly, the MC estimator of the pdf of the lognormal sum has the form

$$\widehat{f}_n(nx) = \mathcal{L}^{n-1}(\widetilde{\theta}(x)) \left[\frac{1}{n}\sum_{i=1}^{n} \mathrm{e}^{\widetilde{\theta}(x)S_{n,-i}}\, f(nx - S_{n,-i}).\right].$$

In our numerical results we have used $R = 100,000$. Tables 5-9 contain the numerical results for the CDF of the lognormal sum for various combinations of the parameters σ and n. Results for the PDF are given in Tables 10-11. The approximations become sharper as either σ or x tend to 0.

Table 5: Approximation of the CDF of a lognormal sum with $n = 4$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MC/MC
0.6500	2.60	11.38	1.53e-04	1.61e-04 ± 1.83e-06
0.7000	2.80	8.87	1.24e-03	1.30e-03 ± 1.37e-05
0.7500	3.00	6.81	6.54e-03	6.92e-03 ± 6.66e-05
0.8000	3.20	5.09	2.40e-02	2.55e-02 ± 2.24e-04
0.8500	3.40	3.65	6.59e-02	7.11e-02 ± 5.61e-04
0.9000	3.60	2.43	1.44e-01	1.55e-01 ± 1.10e-03

Table 6: Approximation of the CDF of a lognormal sum with $n = 64$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MC/MC
0.9219	59.00	1.95	2.00e-04	2.04e-04 ± 2.56e-06
0.9336	59.75	1.71	8.31e-04	8.56e-04 ± 1.01e-05
0.9453	60.50	1.48	2.97e-03	3.06e-03 ± 3.35e-05
0.9570	61.25	1.26	8.99e-03	9.22e-03 ± 9.43e-05
0.9688	62.00	1.04	2.40e-02	2.45e-02 ± 2.30e-04
0.9805	62.75	0.83	5.45e-02	5.59e-02 ± 4.79e-04

Table 7: Approximation of the CDF of a lognormal sum with $n = 256$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MC/MC
0.9727	249.00	0.97	9.96e-05	1.06e-04 ± 1.38e-06
0.9805	251.00	0.83	6.47e-04	6.75e-04 ± 8.16e-06
0.9844	252.00	0.76	1.52e-03	1.57e-03 ± 1.82e-05
0.9883	253.00	0.70	3.52e-03	3.43e-03 ± 3.79e-05
0.9922	254.00	0.63	6.70e-03	7.02e-03 ± 7.41e-05
1.0000	256.00	0.50	2.39e-02	2.50e-02 ± 2.37e-04

Table 8: Approximation of the CDF of a lognormal sum with $n = 64$ and $\sigma = 0.125$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MCIS/MC
0.9500	60.80	3.98	8.37e-05	8.51e-05 ± 1.08e-06
0.9563	61.20	3.52	4.15e-04	4.16e-04 ± 4.97e-06
0.9625	61.60	3.06	1.70e-03	1.69e-03 ± 1.90e-05
0.9688	62.00	2.61	5.87e-03	5.88e-03 ± 6.15e-05
0.9750	62.40	2.17	1.73e-02	1.78e-02 ± 1.71e-04
0.9812	62.80	1.74	4.40e-02	4.45e-02 ± 3.94e-04

Table 9: Approximation of the CDF of a lognormal sum with $n = 64$ and $\sigma = 0.072$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MCIS/MC
0.9703	62.10	6.51	1.44e-04	1.43e-04 ± NaN
0.9734	62.30	5.85	5.35e-04	5.37e-04 ± NaN
0.9766	62.50	5.20	1.76e-03	1.78e-03 ± 1.98e-05
0.9797	62.70	4.55	5.17e-03	5.26e-03 ± 5.51e-05
0.9828	62.90	3.91	1.36e-02	1.37e-02 ± 1.34e-04
0.9859	63.10	3.28	3.17e-02	3.21e-02 ± 2.92e-04

Table 10: Approximation of the PDF of a lognormal sum with $n = 4$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MC/MC
0.6500	2.60	11.38	1.90e-03	1.88e-03 ± 8.66e-06
0.7000	2.80	8.87	1.23e-02	1.22e-02 ± 5.62e-05
0.7500	3.00	6.81	5.15e-02	5.08e-02 ± 2.36e-04
0.8000	3.20	5.09	1.49e-01	1.47e-01 ± 6.85e-04
0.8500	3.40	3.65	3.16e-01	3.16e-01 ± 1.47e-03
0.9000	3.60	2.43	5.25e-01	5.24e-01 ± 2.46e-03

Table 11: Approximation of the PDF of a lognormal sum with $n = 64$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MC/MC
0.9219	59.00	1.95	4.19e-04	4.12e-04 ± 5.73e-06
0.9336	59.75	1.71	1.55e-03	1.55e-03 ± 2.14e-05
0.9453	60.50	1.48	4.90e-03	4.83e-03 ± 6.72e-05
0.9570	61.25	1.26	1.29e-02	1.28e-02 ± 1.79e-04
0.9688	62.00	1.04	2.98e-02	2.91e-02 ± 4.05e-04
0.9805	62.75	0.83	5.72e-02	5.64e-02 ± 7.85e-04

Table 12: Approximation of the PDF of a lognormal sum with $n = 256$ and $\sigma = 0.25$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MCIS/MC
0.9727	249.00	0.97	1.03e-04	1.05e-04 ± 2.15e-06
0.9805	251.00	0.83	5.84e-04	5.86e-04 ± 1.21e-05
0.9844	252.00	0.76	1.27e-03	1.26e-03 ± 2.62e-05
0.9883	253.00	0.70	2.73e-03	2.57e-03 ± 5.30e-05
0.9922	254.00	0.63	4.78e-03	4.87e-03 ± 1.00e-04
1.0000	256.00	0.50	1.42e-02	1.45e-02 ± 2.98e-04

Table 13: Approximation of the PDF of a lognormal sum with $n = 64$ and $\sigma = 0.125$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MCIS/MC
0.9500	60.80	3.98	3.55e-04	3.58e-04 ± 4.82e-06
0.9563	61.20	3.52	1.58e-03	1.56e-03 ± 2.11e-05
0.9625	61.60	3.06	5.73e-03	5.70e-03 ± 7.70e-05
0.9688	62.00	2.61	1.73e-02	1.72e-02 ± 2.34e-04
0.9750	62.40	2.17	4.41e-02	4.49e-02 ± 6.05e-04
0.9812	62.80	1.74	9.54e-02	9.56e-02 ± 1.30e-03

Table 14: Approximation of the PDF of a lognormal sum with $n = 64$ and $\sigma = 0.072$.

x	nx	$\widetilde{\theta}(x)$	Saddle/MC	MCIS/MC
0.9703	62.10	6.51	9.99e-04	9.84e-04 ± NaN
0.9734	62.30	5.85	3.39e-03	3.39e-03 ± NaN
0.9766	62.50	5.20	1.01e-02	1.01e-02 ± 1.36e-04
0.9797	62.70	4.55	2.65e-02	2.68e-02 ± 3.59e-04
0.9828	62.90	3.91	6.19e-02	6.17e-02 ± 8.29e-04
0.9859	63.10	3.28	1.26e-01	1.26e-01 ± 1.69e-03

References

[1] *Basel II: Revised International Capital Framework*. Bank for International Settlements, 2004.

[2] I. Aitchison and J. A. C. Brown. *The lognormal distribution with special reference to its uses in economics*. Cambridge University Press, Cambridge, 1957.

[3] S. Asmussen. *Applied Probability and Queues*. Springer-Verlag, New York, 2003.

[4] S. Asmussen, K. Binswanger, and B. Højgaard. Rare events simulation for heavy-tailed distributions. *Bernoulli*, 6:303–322, 2000.

[5] S. Asmussen, J. Blanchet, S. Juneja, and L. Rojas-Nandayapa. Efficient simulation of tail probabilities of sums of correlated lognormals. *Annals of Operations Research*, 189:5–23, 2011.

[6] S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Springer-Verlag, New York, 2007.

[7] S. Asmussen, J. L. Jensen, and L. Rojas-Nandayapa. On the Laplace transform of the lognormal distribution. *Submitted*, 2014.

[8] S. Asmussen and L. Rojas-Nandayapa. Asymptotics of sums of lognormal random variables with Gaussian copula. *Stat. Probab. Lett.*, 78, 2008. To appear.

[9] A. Bassamboo, S. Juneja, and A. Zeevi. Portfolio credit risk with extremal dependence: Asymptotic analysis and efficient simulation. *Operations Research*, 56:593–606, 2008.

[10] J. H. Blanchet and L. Rojas-Nandayapa. Efficient simulation of tail probabilities of sums of dependent random variables. *Journal of Applied Probability*, 48A:147–164, 2011.

[11] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5:329–359, 1996.

[12] E. L. Crow and K. Shimizu. *Lognormal distributions: theory and applications*. Marcel Dekker Inc., New York, 1988.

[13] D. Dufresne. The log-normal approximation in financial and other computations. *Advances in Applied Probability*, pages 747–773, 2004.

[14] D. Dufresne. Sums of lognormals. Technical report, Centre for Actuarial Sciences, University of Melbourne, 2009.

[15] L. Fenton. The sum of log-normal probability distributions in scatter transmission systems. *IRE Transactions on Communications Systems*, 8(1):57–67, 1960.

[16] S. Foss and A. Richards. On sums of conditionally independent subexponential random variables. *Mathematics of Operations Research*, 35:102–119, 2010.

[17] J. A. Gubner. A new formula for lognormal characteristic functions. *IEEE Transactions on Vehicular Technology*, 55(5):1668–1671, 2006.

[18] A. Gulisashvili and P. Tankov. Tail behavior of sums and differences of log-normal random variables. *arXiv:1309.3057*, 2014.

[19] J. L. Jensen. *Saddlepoint Approximations*. Oxford Science Publications, Oxford, 1994.

[20] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 1. Wiley, New York, 2 edition, 1994.

[21] E. Limpert, W. A. Stahel, and M. Abbt. Log-normal distributions across the sciences: Keys and clues. *BioScience*, 51:341–352, 2001.

[22] N. A. Marlow. A normal limit theorem for power sums of normal random variables. *Bell System Technical Journal*, 46:2081–2089, 1967.

[23] M. A. Milevsky and S. E. Posner. Asian options, the sum of lognormals, and the reciprocal gamma distribution. *The Journal of Financial and Quantitative Analysis*, 1998.

[24] A. Mitra and S. I. Resnick. Aggregation of risks and asymptotic independence. *Advances in Applied Probability*, pages 797–828, 2009.

[25] L. Rojas-Nandayapa. *Risk probabilities: Asymptotics and Simulation*. PhD thesis, Aarhus University, 2008.

[26] O. Thorin and N. Wikstad. Calculation of ruin probabilities when the claim distribution is lognormal. *Astin Bulletin*, pages 231–246, 1977.

A Appendix: R.V. generation

Let us consider the A-R algorithm 4.2 for generating from the exponential family generated by the lognormal. Here the rejection probability is given by $\mathbb{E}[\mathrm{e}^{-\log^2(Z)/2\sigma^2}]$ where $Z \sim \text{Gamma}(\alpha+1,\alpha)$ and $\alpha = \mathcal{W}(\theta\sigma^2\mathrm{e}^{-\sigma^2})/\sigma^2$. After some manipulations we can rewrite the acceptance probability as

$$\begin{aligned}
\mathbb{E}[\mathrm{e}^{-\log^2(Z)/2\sigma^2}] &= \int_0^\infty \mathrm{e}^{-\log^2 z/2\sigma^2}\frac{\alpha^{\alpha+1}z^\alpha\mathrm{e}^{-\alpha z}}{\Gamma(\alpha+1)}\,\mathrm{d}z \\
&= \frac{\alpha^{\alpha+1}}{\Gamma(\alpha+1)}\sqrt{2\pi}\sigma\int_0^\infty \frac{1}{\sqrt{2\pi}\sigma z}\mathrm{e}^{-(\log z-\sigma^2(\alpha+1))^2/2\sigma^2+\sigma^2(\alpha+1)^2/2-\alpha z}\,\mathrm{d}z \\
&= \frac{\alpha^{\alpha+1}}{\Gamma(\alpha+1)}\sqrt{2\pi}\sigma\,\exp\left\{\frac{\sigma^2(\alpha+1)^2}{2}\right\}\mathcal{L}\Big(\alpha\mathrm{e}^{\sigma^2(\alpha+1)}\Big)
\end{aligned}$$

where $\mathcal{L}(\cdot)$ is the Laplace transform of the lognormal distribution. From the previous expression it can be observed that when $\theta \to 0$ then $\alpha \to 0$ and also the probability of acceptance goes to 0.

In contrast, notice that if $\theta \to 0$ then the probability of acceptance of algorithm 4.1 goes to 1. Thus it seems natural to combine algorithms 4.1 and 4.2 into a single one. We have expressed the probabilities of acceptance as functions of the Laplace transform of the lognormal distribution so we can insert our approximations, compare their values an choose the algorithm with the largest approximation of the probability of acceptance.

B Appendix: On the estimation of $\mathcal{L}^n(\theta)$.

We first recall the estimator of $\mathcal{L}(\theta)$ suggested in [7]:

$$\widehat{\mathcal{L}}(\theta) = \widetilde{\mathcal{L}}(\theta)\,\vartheta(Y,\theta), \tag{B.1}$$

where

$$\widetilde{\mathcal{L}}(\theta) = \exp\Big\{-\frac{\mathcal{W}^2(\theta\sigma^2)+2\,\mathcal{W}(\theta\sigma^2)}{2\sigma^2}\Big\}, \qquad \vartheta(t;\theta) = \exp\Big\{-\frac{\mathcal{W}(\theta\sigma^2)}{\sigma^2}\big(\mathrm{e}^t-1-t\big)\Big\},$$

and $Y \sim \mathrm{N}(0,\sigma^2)$. This was proved there to be unbiased and logarithmically efficient.

Remark B.1. If σ^2 is small or modest (as in our numerical examples), $\mathrm{e}^Y - 1 - Y$ is not too far from Y^2. This suggests using

$$\widetilde{\mathcal{L}}(\theta)\exp\Big\{-\frac{\mathcal{W}(\theta\sigma^2)}{\sigma^2}Y^2\Big\} \tag{B.2}$$

as control variate ([6, V.2]) for $\widehat{\mathcal{L}}(\theta)$. In fact, the requirement for this method that the mean is known is satisfied because the expectation of (B.2) comes out as $\big(\mathcal{W}(\theta\sigma^2)+1\big)^{1/2}$. We have not implemented this control variate idea.

Now consider the problem of estimating $\mathcal{L}^n(\theta)$. Let $\widehat{\ell}_R$ be (B.1) averaged over R replications. so that $\widehat{\ell}_R - \mathcal{L}(\theta)$ is approximately normal$(0,\tau^2/R)$ for some τ^2 that can be estimated by the empirical variance t^2. For any smooth function φ, we can write

$$\varphi(\widehat{\ell}_R) \;=\; \varphi\big(\mathcal{L}(\theta)\big) + \varphi'\big(\mathcal{L}(\theta)\big)\big(\widehat{\ell}_R - \mathcal{L}(\theta)\big) + \frac{\varphi''\big(\mathcal{L}(\theta)\big)}{2}\big(\widehat{\ell}_R - \mathcal{L}(\theta)\big)^2 + O(R^{-3/2})\,.$$

Taking $\varphi(v) = v^n$, this leads to two observations:
a) using the estimator

$$\widehat{\ell}_R - \frac{n(n-1)\widehat{\ell}_R^{\,n-2}}{2R}t^2$$

instead of $\widehat{\ell}_R^{\,n}$ reduces the bias from order R^{-1} to order $R^{-3/2}$;
b) the variance of the (biased) estimator $\widehat{\ell}_R^{\,n}$ can be estimated by $n^2\widehat{\ell}_R^{\,2n-2}t^2/R$
[these arguments are essentially the delta method, [6, III.3]].

An alternative to the estimator $\widehat{\ell}^n$ is to take the average over R replications over the product of n independent copies of $\widehat{\mathcal{L}}(\theta)$. Clearly this estimator is unbiased but of course more costly to produce. The variance can just be estimated by standard Monte Carlo.

C Appendix: Proof of Proposition 4.7

We will need the following Lemma:

Lemma C.1. *Let*

$$\widetilde{\beta}_n(u) := \mathrm{e}^{\widetilde{\theta}(u)S_n}\left[\mathcal{L}(\widetilde{\theta}(u))\right]^n \mathbb{I}\{S_n < nu\}. \tag{C.1}$$

Then

$$\limsup_{u\to 0} \frac{\mathrm{MSE}_{\alpha_n(u)}[\widehat{\alpha}_n(u)]}{\mathbb{E}_{\widetilde{\theta}(u)}\left[\widehat{\beta}_n^2(u)\right](1+\gamma(u))^{n/2}} \leq 2.$$

Proof of Lemma C.1. Note that $\widehat{\beta}_n(u)$ is unbiased. This follows from the change of measure argument

$$\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n(u)] = \mathbb{E}_{\widetilde{\theta}(u)}\left[\,\mathrm{e}^{\widetilde{\theta}(u)S_n+n\kappa(\widetilde{\theta}_u)}\,; S_n < nu\right] = \mathbb{E}[\mathbb{I}(S_n < nu)] = \mathbb{P}(S_n < nu).$$

Hence we can write

$$\limsup_{u\to 0} \frac{\mathrm{MSE}_{\alpha_n(u)}[\widehat{\alpha}_n(u)]}{\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n^2(u)](1+\gamma(u))^{n/2}} = \limsup_{u\to 0} \frac{\mathbb{E}_{\widetilde{\theta}(u)}\left[\left(\widehat{\alpha}_n(u) - \mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n(u)]\right)^2\right]}{\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n^2(u)](1+\gamma(u))^{n/2}}.$$

Since all terms are positive then the last expression can be bounded with

$$\limsup_{u\to 0} \frac{\mathbb{E}_{\widetilde{\theta}(u)}\left[\widehat{\alpha}_n^2(u)\right] + \mathbb{E}^2_{\widetilde{\theta}(u)}\left[\widehat{\beta}_n(u)\right]}{\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n^2(u)](1+\gamma(u))^{n/2}} < \limsup_{u\to 0} \frac{\mathbb{E}_{\widetilde{\theta}(u)}\left[\widehat{\alpha}_n^2(u)\right] + \mathbb{E}_{\widetilde{\theta}(u)}\left[\widehat{\beta}_n^2(u)\right]}{\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n^2(u)](1+\gamma(u))^{n/2}}.$$

Using the definition for both $\widehat{\alpha}_n(u)$ and $\widehat{\beta}_n(u)$ we arrive at

$$\limsup_{u\to 0} \frac{\mathbb{E}_{\widetilde{\theta}(u)}\left[\mathrm{e}^{-2\widetilde{\theta}(u)S_n}; S_n < nu\right]\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\mathcal{L}}^{2n}_{IS}(\widetilde{\theta}(u))] + \mathbb{E}_{\widetilde{\theta}(u)}\left[\mathrm{e}^{-2\widetilde{\theta}(u)S_n}; S_n < nu\right]\mathcal{L}^{2n}(\widetilde{\theta}(u))}{\mathbb{E}_{\widetilde{\theta}(u)}\left[\mathrm{e}^{-2\widetilde{\theta}(u)S_n}; S_n < nu\right]\mathcal{L}^{2n}(\widetilde{\theta}(u))\,(1+\gamma(u))^{n/2}}.$$

The last simplifies as

$$\limsup_{u\to 0} \frac{\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\mathcal{L}}^{2n}_{IS}(\widetilde{\theta}(u))] + \mathcal{L}^{2n}(\widetilde{\theta}(u))}{\mathcal{L}^{2n}(\widetilde{\theta}(u))\,(1+\gamma(u))^{n/2}}.$$

Now, since $\widehat{\mathcal{L}}_{\mathrm{IS}}(\cdot)$ is unbiased, then the following upper bound is obtained by a direct application of Jensen's inequality

$$\limsup_{u\to 0} \frac{2\,\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\mathcal{L}}^{2n}_{IS}(\widetilde{\theta}(u))]}{\mathcal{L}^{2n}(\widetilde{\theta}(u))(1+\gamma(u))^{n/2}}.$$

By using the formula (4.4) and recalling that $\mathcal{W}(\widetilde{\theta}(u)\sigma^2) = \gamma(u)$ we arrive to the equivalent expression

$$\limsup_{u\to\infty} \frac{2\,\widetilde{\mathcal{L}}^{2n}(\widetilde{\theta}(u))\mathbb{E}[\vartheta^{2n}_k(Z;\widetilde{\theta}(u))]}{\widetilde{\mathcal{L}}^{2n}(\widetilde{\theta}(u))} \leq 2.$$

The last inequality follows from the fact that $\vartheta(\cdot) \leq 1$. □

Proof of Proposition 4.7. By Lemma C.1 we have that

$$\limsup_{u\to 0}\frac{\mathrm{MSE}_{\alpha_n(u)}[\widehat{\alpha}_n(u)]}{\alpha_n^{2-\epsilon}(u)} \le \limsup_{u\to 0}\frac{2\,\mathbb{E}_{\widetilde{\theta}(u)}[\widehat{\beta}_n^2(u)](1+\gamma(u))^{n/2}}{\alpha_n^{2-\epsilon}(u)}.$$

Using the change of measure argument we can rewrite the expectation in the numerator in the expression above as

$$\begin{aligned}\mathbb{E}_{\theta(u)}[\widehat{\beta}_n^2(u)] &= \mathbb{E}_{\widetilde{\theta}(u)}\Big[\,\mathrm{e}^{2\widetilde{\theta}(u)S_n+2n\kappa(\widetilde{\theta}(u))}\,;S_n<nu\Big]\\ &= \mathbb{E}\Big[\mathrm{e}^{\widetilde{\theta}(u)S_n+n\kappa(\widetilde{\theta}(u))};S_n<nu\Big]. \qquad \text{(C.2)}\end{aligned}$$

Now, observe that $\widetilde{\theta}(u)S_n<\widetilde{\theta}(u)nu$ in the set $\{S_n<nu\}$, so we obtain the bound

$$\mathrm{e}^{\widetilde{\theta}(u)nu+n\kappa(\widetilde{\theta}(u))}\,\mathbb{E}[\mathbb{I}(S_n<nu)] = \mathrm{e}^{\widetilde{\theta}(u)nu-n\kappa(\widetilde{\theta}(u))}\mathbb{P}(S_n<nu).$$

Also we have the following inequalities

$$\alpha(u)=\mathbb{P}(S_n<nu)\ge\mathbb{P}(\max\{X_i\}<u)=\mathbb{P}^n(Y_1<u).$$

Putting together these results we conclude that

$$\begin{aligned}\lim_{u\to 0}\frac{\mathbb{E}_\theta[\widehat{\beta}^2(u)](1+\gamma(u))^{n/2}}{\alpha^{2-\epsilon}(u)} &< \lim_{u\to 0}\frac{\mathrm{e}^{\widetilde{\theta}(u)nu+n\kappa(\widetilde{\theta}(u))}(1+\gamma(u))^{n/2}}{\alpha^{1-\epsilon}(S_n<nu)}\\ &< \lim_{u\to 0}\left[\frac{\mathrm{e}^{\widetilde{\theta}(u)u+\kappa(\widetilde{\theta}(u))}\sqrt{1+\gamma(u)}}{\mathbb{P}^{1-\epsilon}(Y_1<u)}\right]^n\\ &= \left[\lim_{u\to 0}\frac{\mathrm{e}^{\widetilde{\theta}(u)u+\kappa(\widetilde{\theta}(u))}\sqrt{1+\gamma(u)}}{\mathbb{P}^{1-\epsilon}(Y_1<u)}\right]^n,\end{aligned}$$

where the last equality holds provided that the limit inside the brackets exists. In fact, we will prove that such limit is 0. For the numerator in the limit above we use (1.2) to obtain

$$\mathrm{e}^{\widetilde{\theta}(u)u+\kappa(\widetilde{\theta}(u))}\sqrt{1+\gamma(u)}\sim\exp\left\{\widetilde{\theta}(u)u-\frac{\gamma^2(u)+2\gamma(u)}{2\sigma^2}\right\}\sqrt{1+\gamma(u)},$$

while for the denominator we employ *Mill's ratio* so we get

$$\mathbb{P}(Y_1<u)\sim\frac{\sigma}{\sqrt{2\pi}|\log u|}\exp\left\{-\frac{\log^2 u}{2\sigma^2}\right\},\qquad u\to 0.$$

Hence

$$\begin{aligned}&\lim_{u\to 0}\frac{\mathrm{e}^{\widetilde{\theta}(u)u+\kappa(u)}}{\mathbb{P}^{1-\epsilon}(Y_1<u)}\\ &\quad=\lim_{u\to 0}\frac{c_1\sqrt{1+\gamma(u)}}{|\log u|^{-(1-\epsilon)}}\exp\left\{-\frac{\gamma^2(u)+2\gamma(u)}{2\sigma^2}+u\widetilde{\theta}(u)+(1-\epsilon)\frac{\log^2 u}{2\sigma^2}\right\}\\ &\quad=\lim_{u\to 0}\frac{c_1\sqrt{1+\gamma(u)}}{|\log u|^{-(1-\epsilon)}}\exp\left\{-\frac{\gamma(u)}{\sigma^2}+u\widetilde{\theta}(u)-\frac{\gamma^2(u)-(1-\epsilon)\log^2 u}{2\sigma^2}\right\},\end{aligned}$$

for some constant c_1. We apply Lemmas C.3–C.4 to prove that the last limit is

$$\lim_{u\to 0} c_2 \,|\log u|^{3/2-\epsilon} \exp\left\{ -\epsilon\frac{\log^2 u}{2\sigma^2} + O(|\log u|^{-2})\right\}.$$

The last limit is equal to 0 for all $\epsilon > 0$. This completes the proof. □

Lemma C.2.

$$\gamma(u) = |\log(u)| + \frac{\sigma^2}{|\log u|} + O(|\log u|^{-2}), \qquad u \to 0.$$

Proof.

$$\begin{aligned}
\gamma(u) &= \frac{-1 + |\log u| + \big(1 + |\log u|\big)\sqrt{1 + \dfrac{2\sigma^2}{(1+|\log u|)^2}}}{2} \\
&= \frac{-1 + |\log u| + \big(1 + |\log u|\big)\big(1 + \frac{\sigma^2}{\log^2 u} + O(\log^{-3} u)\big)}{2} \\
&= |\log u| + \frac{\sigma^2}{|\log u|} + O(|\log u|^{-2}).
\end{aligned}$$

□

Lemma C.3.

$$u\widetilde{\theta}(u) = \frac{|\log u|}{\sigma^2} + 1 + O(|\log u|^{-1}).$$

Proof. From Lemma C.2 it follows that for $u > 0$ close enough to 0 it holds that

$$u\mathrm{e}^{\gamma(u)} = \exp\left\{\frac{\sigma^2}{|\log u|} + O(|\log u|^{-2})\right\} = 1 + \frac{\sigma^2}{|\log u|} + O(|\log u|^{-2}).$$

Using this result and combining again with Lemma C.2 we get

$$u\widetilde{\theta}(u) = \frac{u\gamma(u)\mathrm{e}^{\gamma(u)}}{\sigma^2} = \frac{|\log u|}{\sigma^2} + 1 + O(|\log u|^{-1})$$

□

Lemma C.4.

$$\frac{\gamma^2(u) - \log^2 u}{2\sigma^2} = 1 + O(|\log u|^{-2}).$$

Proof of Lemma C.4. From Lemma C.2 it follows that

$$\gamma^2(u) - |\log u|^2 = 2\sigma^2 + O(|\log u|^{-2})$$

□